

sub
A2

DYNAMIC ALLOCATION OF SHARED NETWORK RESOURCES
BETWEEN CONNECTION-ORIENTED AND CONNECTIONLESS
TRAFFIC

CROSS-REFERENCE TO RELATED APPLICATIONS

- 5 This is the first application filed for the present invention.

MICROFICHE APPENDIX

Not Applicable.

TECHNICAL FIELD

- 10 The present invention relates to resource management in communications networks, and in particular to dynamic allocation of shared network resources between connection-oriented and connectionless traffic in a communications network.

15 BACKGROUND OF THE INVENTION

- In the modern network space, packetized data traffic of various different protocols (e.g. internet protocol, frame relay, asynchronous transfer mode, etc.) is transported over a common network infrastructure. Each protocol provides its own packet (or frame) size and format standards. Additionally, some protocols (e.g. IP) are specifically designed to allow packets having widely varying lengths. New routing protocols, for example the multi-protocol label switching (MPLS) protocol have been proposed to facilitate multi-protocol traffic across a common network infrastructure.
- 20
- 25

Under the MPLS protocol, label switched packets (LSPs) are propagated across the network hop-by-hop along a

00527504-034700

path that is set up at the beginning of a communications session. In general, the label assigned to the LSP can be different for each hop, with the label conversion being performed by the node serving the respective hop. Resources
5 of each hop (i.e. the node serving the hop) of the path are reserved during set-up of the path, and normally will not be available for carrying other traffic until the path is released.

10 The mapping of an end-to-end path at the beginning of a communications session characterizes the MPLS protocol as "connection oriented". Other protocols, (such as IP, ATM and Frame Relay) which do not transport data over predefined end-to-end paths are referred to as "connectionless". Typically, connectionless traffic is
15 routed across a network using a shortest-path or least-cost-path routing protocol, such as, for example, the Interior Gateway Protocol (IGP). In general, a metric (e.g. a link distance vector, or a link cost factor) is assigned to each link and used within each router for
20 mapping packet destination addresses to downstream links. The metric is normally provisioned for traffic engineering, and reflects not only geographic distances, but also provisioned bandwidth of each link. A higher metric on a particular link makes that link less attractive for
25 carrying connectionless traffic, so that the IGP will normally operate to route connectionless traffic away from that link.

Both connection-oriented and connectionless traffic may be carried over shared network infrastructure. This
30 situation is normally accommodated by adjusting the provisioned IGP metric to reflect an average anticipated amount of bandwidth allocated to the connection-oriented

002760 14922560

Accordingly, during periods of heavy demand for connection-oriented traffic, the provisioned IGP metric for a link may provide an inflated indication of the amount of band-width actually available for connectionless traffic. This can easily result in undesirable congestion on the link. Conversely, during periods of low demand for connection-oriented traffic, the provisioned IGP metric for a link may provide a deflated indication of the amount of band-width actually available for connectionless traffic. This can result in undesirable under-utilization of the link.

SUMMARY OF THE INVENTION

A further object of the present invention is to provide a method of allocating resources between connection oriented and connectionless traffic, by adjusting an IGP metric in accordance with MPLS resource reservations.

Accordingly, an aspect of the present invention provides a shared network element operative within a communications network capable of end-to-end transport of

connection-oriented traffic and connectionless traffic through the shared network element. The shared network element comprises: means for determining a resource requirement of the connection-oriented traffic; and means for dynamically adjusting a respective connectionless traffic metric based on the determined resource requirement of the connection-oriented traffic.

Another aspect of the present invention provides a method of managing an allocation of resources between connection-oriented traffic and connectionless traffic being routed through a shared physical network element of a communications network. The method comprises the steps of: determining a resource requirement of the connection-oriented traffic; and dynamically adjusting a respective connectionless traffic metric based on the determined resource requirement of the connection-oriented traffic.

In embodiments of the invention, the connection-oriented traffic is multi-protocol label switched (MPLS) traffic. In such cases, the step of determining the resource requirement of the connection-oriented traffic comprises the steps of: receiving MPLS reservation requests in respect of the shared physical network element; and dynamically adjusting a total amount of resources required to satisfy the received MPLS reservation requests.

In embodiments of the invention, the connectionless traffic includes internet protocol (IP) packet traffic. In such cases, routing of the connectionless traffic may be controlled using an interior gateway protocol (IGP) routing system adapted to calculate a shortest path route of the connectionless traffic through the communications network,

the shortest path routing being based on a respective metric concerning each physical network element forming the network.

5 The step of dynamically adjusting the respective metric preferably comprises the steps of: increasing the respective metric as the determined resource requirement of the connection-oriented traffic increases; and decreasing the respective metric as the determined resource requirement of the connection-oriented traffic decreases.

10 In some embodiments of the invention, the respective metric may be a link distance vector associated with a respective link connected to a node of the communications network. In such cases, the step of dynamically adjusting the respective metric comprises the
15 steps of: determining an updated value of the link distance vector; and updating a mapping table maintained by the node with the updated value of the link distance vector.

An updated value of the link distance vector may be determined by querying a resource allocation table
20 comprising a plurality of characteristic resource allocation values and a respective link distance vector value corresponding to each characteristic resource allocation value. Querying the resource allocation table may include the steps of: identifying the characteristic
25 resource allocation value which most closely matches the determined resource requirement of the connection-oriented traffic; and selecting the corresponding link distance vector as the updated link cost factor.

30 In other embodiments of the invention, the respective metric may be a link cost factor associated with a respective link connected to a node of the communications

0027504 03400

network. In such cases, the step of dynamically adjusting the respective metric comprises the steps of: determining an updated value of the link cost factor; updating a PATH database maintained by the node with the updated link cost factor value; and propagating a link state packet containing the updated link cost factor value to neighboring nodes within the network.

An updated value of the link cost factor can be determined by querying a resource allocation table comprising a plurality of characteristic resource allocation values and a respective link cost factor value corresponding to each characteristic resource allocation value. Querying the resource allocation table may comprise the steps of: identifying the characteristic resource allocation value which most closely matches the determined resource requirement of the connection-oriented traffic; and selecting the corresponding link cost factor as the updated link cost factor.

An advantage of the present invention is that by adjusting the IGP metric for a link in accordance with changing resource requirements of connection oriented traffic, routing of connectionless traffic is automatically altered to make effective use of remaining resources while avoiding congestion. No modification of conventional (e.g. IGP) routing methodologies are required to accomplish this result.

BRIEF DESCRIPTION OF THE DRAWINGS

Further features and advantages of the present invention will become apparent from the following detailed description, taken in combination with the appended drawings, in which:

Fig. 1. is a block diagram illustrating a communications network usable in conjunction with an embodiment of the present invention; and

Fig. 2 is a flow chart illustrating exemplary steps in a process for managing resource allocation in accordance with an embodiment of the present invention.

It will be noted that throughout the appended drawings, like features are identified by like reference numerals.

10 **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT**

The present invention provides a system for allocating resources between connectionless and connection-oriented traffic through a communications network. As shown in Fig. 1, a communications network 2 usable in conjunction with the present invention generally comprises a plurality of routers 4 (four are shown in Fig. 1) interconnected by links 6. The links 6 may be fiber optic links. The routers 4 may be agile or non-agile optical routers, and may be configured for wave division multiplex (WDM) and/or dense wave division multiplex (DWDM) transport of packet data traffic. Communications devices 8, for example end user personal computers (PCs) or local area network (LAN) servers may be connected to the communications network 2 via one or more access points 10. The communications network 2 may also be connected to one or more federated networks 12, for example an ATM or an IP network, through a respective gateway 14.

In the example of Fig. 1, connection-oriented traffic is conveyed through an end-to-end MPLS path 16 mapped across the communications network 2 between a source

node and a destination node via one or more intervening routers 4. The path 16 is divided into hops 18, each of which is served a respective node (e.g. the source node or a router 4) connected at the up-stream end of the respective hop 18. In the illustrated example, the source and destination nodes are located at respective access points 10a and 10b, and two intervening routers 4a and 4b are incorporated into the path 16. A first hop 18a of the path 16 extends between the source node at access point 10a and a first router 4a. A second hop 18b extends between the routers 4a and 4b. Finally, a third hop 18c extends between the router 4b and the destination node at access point 10b.

In addition, connectionless traffic is transported through the communications network between the gateway 14 and the destination node 10b. Routing of the connectionless traffic is handled, for example, in accordance with conventional least cost path routing using the Interior Gateway Protocol (IGP) to map destination addresses to downstream links. In the example illustrated in Fig. 1, a least cost path 20 calculated using provisioned IGP metrics for each link follows the route indicated by dashed arrows. As may be seen in Fig. 1, this least cost path 20 shares two hops (18b and 18c) with the MPLS path 16, which may lead to congestion on those hops. Thus the present invention provides a technique of managing the allocation of resources between the two traffic flows, in order to avoid congestion within the shared hops.

In accordance with the present invention, congestion within shared physical elements of the network (e.g. within the routers 4a,4b and links 6 of the shared hops 18b,18c) is avoided by adjusting the IGP metrics concerning each of the shared links 6 to reflect the

resources (e.g. bandwidth) which have been allocated to the MPLS path 16. Adjustment of the IGP metric in this manner makes each of the shared links 6 less attractive for either shortest-distance-path routing or least-cost-path routing.

5 In cases where either the bandwidth allocated to the MPLS path 16, or the volume of connectionless traffic is large, the adjusted IGP metrics on the shared links 6 may cause the routing protocol to favour an alternative path 22 (indicated by dotted arrows in Fig. 1) which avoids sharing
10 links 6 with the MPLS path 16.

Adjustment of IGP metrics can be based on a resource allocation table 24 which may be co-resident within each router 4 or centrally located and accessible by each router 4 through the network 2. The resource
15 allocation table 24 generally operates to receive a query from a node (any of access points 10, routers 4, or gateway 14) containing a value indicative of resources allocated to connection oriented traffic (in the present example bandwidth allocated to the MPLS path 16). The resource
20 allocation table 24 returns a response message to the node containing an updated value of an IGP metric. The node can then update its local path database (not shown) and propagate a link state packet to neighbouring nodes of the network so that conventional IGP routing can proceed on the
25 basis of the updated IGP metric. Exemplary steps in a process of adjusting the IGP metric for a link 6 connected to a router 4 within the network 2 are described below with respect to Fig. 2.

Adjustment of IGP metrics for each link 6 is
30 preferably performed by the associated node during the set up of an MPLS path 16 through that node. Accordingly, as shown in Fig. 2, when an MPLS reservation request is

002750 034000

received at a node (step 100), the request is processed in a conventional manner (step 102) and resources of the node (and a downstream link) are allocated to a hop of an MPLS path 16 being set up across the communications network 2.

5 Once resources have been allocated to the MPLS path 16, the node operates to determine a total resource allocation to connection oriented traffic over the involved downstream link 6 (step 104). The node then formulates a query message containing the total amount of allocated resources (step

10 106) and forwards the query message to the resource allocation table 24. Upon receipt of the query message from the node, the resource allocation table 24 determines an adjusted IGP metric value (step 108) and sends a response message to the node containing the adjusted IGP metric

15 value. If the response message is received by the node prior to expiry of a predetermined time out period (step 110), the node operates to update its local path database with the adjusted IGP metric value for the downstream link involved in the newly set up MPLS path 16 (step 112). This

20 change in the local path database of the node automatically triggers the generation of a link state packet which is propagated to neighbouring nodes within the network for handling in a conventional manner (step 114).

While not shown in the drawings, it will be

25 appreciated that an analogous process may be executed within a node to adjust the IGP metric as resources allocated to connection oriented traffic are released (e.g. as an MPLS path is taken down). Thus when an MPLS path release message is received at a node, the message is

30 processed in a conventional manner, and resources of the node (and a downstream link) allocated to a hop of an MPLS path across the communications network are released. Once resources have been released, the node operates to

00527504-034700

determine a total amount of resources allocated to connection oriented traffic over the involved downstream link. The node then formulates a query message containing the total resource allocation, and forwards the query message to the resource allocation table. Upon receipt of the query message from the node, the resource allocation table determines an adjusted IGP metric value and sends a response message to the node containing the adjusted IGP metric value. If the response message is received by the node prior to a predetermined time out period, the node operates to update its local path database with the adjusted IGP metric value for the downstream link involved in the recently released MPLS path. This change in the local path database of the node automatically triggers the generation of a link state packet which is propagated to neighbouring nodes within the network for handling in a normal manner

Various techniques may be used within the resource allocation table 24 to determine adjusted values for the IGP metric. For example, the resource allocation table 24 may contain a list of characteristic resource allocation values, and corresponding IGP metric values. In this case, upon receipt of a query message from a node, the resource allocation is extracted from the query message and compared to the characteristic values in the resource allocation table 24. The IGP metric value corresponding to the characteristic value which most closely matches the resource allocation extracted from the query message is then selected as the adjusted IGP metric, and sent back to the node in the response message. In an alternative embodiment, and in order to accommodate differing bandwidth capacities of links within the network, at least the characteristic resource allocation values within the

resource allocation table 24 may be proportional, and so may represent a fraction of the total available bandwidth that has been allocated to connection-oriented traffic.

The embodiment(s) of the invention described above
5 is(are) intended to be exemplary only. The scope of the
invention is therefore intended to be limited solely by the
scope of the appended claims.